

MACHINE LEARNING

what it is and how to get started

Auralee Edelen
May 4, 2017

MACHINE LEARNING

what it is and how to get started


Auralee Edelen
May 4, 2017

Caveats!

- *20 minutes → very high-level overview*
- *heavy bias toward neural networks*



Field Taxonomy (as of now...)

- Artificial Intelligence (AI)
 - *Field of getting machines to exhibit aspects of human intelligence, esp. knowledge, learning, planning, reasoning, perception*
 - Narrow AI: focused on a task or similar set of tasks
 - General AI: human-equivalent or greater performance on any task
- Machine Learning (ML)
 - *Field of getting machines to complete tasks without being explicitly programmed*
 - Common tasks: Regression, Classification, Clustering, Dimensionality Reduction
- Neural Networks (NNs)
 - *A set of tools within ML that uses a many connected processing units*
 - Many kinds: feed-forward, recurrent, adversarial, self-organizing maps
 - Very popular right now (somewhere at the top of the hype cycle...)
- Deep Learning (DL)
 - *Learning hierarchical representations*
 - Right now, largely synonymous with methods based on deep (many-layered) NNs



Note that these definitions are not rigid: there is a lot of fluidity in the field at the moment!


Basic Learning Paradigms



Example: Regression using a NN

Data set of **input** and **output** pairs:


$$\left. \begin{matrix} x_1 \\ \vdots \\ x_n \end{matrix} \right\} \begin{matrix} \mathbf{x}_{t_1} \\ \vdots \\ \mathbf{x}_{t_n} \end{matrix} \quad \begin{matrix} \mathbf{y}_{t_1} \\ \vdots \\ \mathbf{y}_{t_n} \end{matrix}$$



Want to find approximate map:

$$\mathbf{g}(\mathbf{x}) = \mathbf{y}$$

Model Learning




Example:

$$C(w, b) = \frac{1}{2t_n} \left[\sum_{t_n} (y_{t_n} - \tilde{y}_{t_n})^2 \right]$$

$$w_k \rightarrow w'_k = w_k - \alpha \frac{\partial C}{\partial w_k}$$


$$b_k \rightarrow b'_k = b_k - \alpha \frac{\partial C}{\partial b_k}$$

ANN Basic Structure

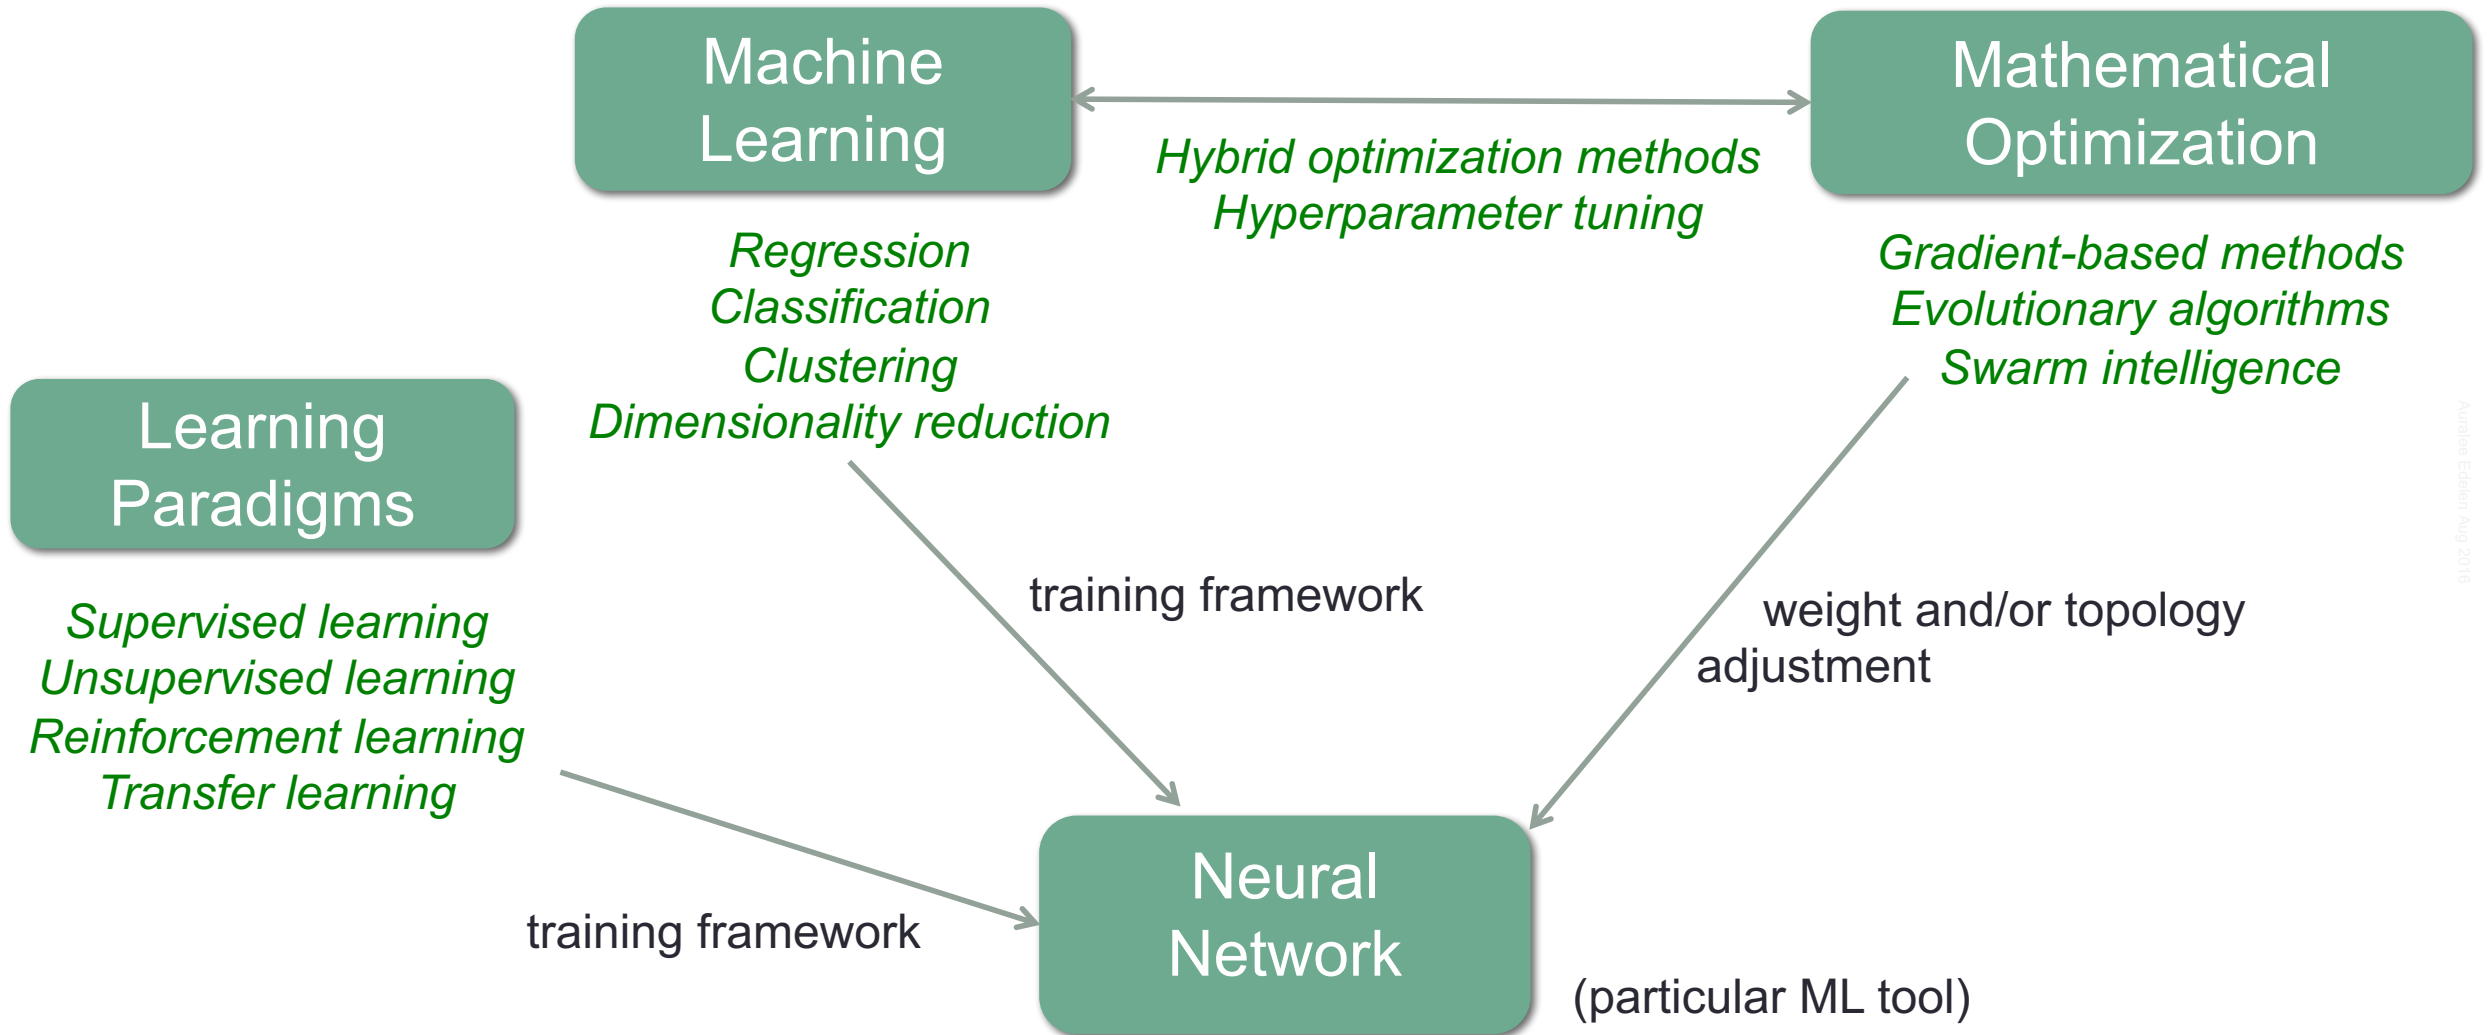


$$f \left(\sum_n w_n x_n + b \right) = a$$

e.g. $f(z) = \frac{2}{(1+e^{-2z})} - 1$



Example of how this all fits together for NNs



ML Software and Related Libraries

- **Theano** – library for fast numerical computation (graph-based, automatic differentiation, python)
- **Tensor Flow** – library for fast numerical computation (graph-based, automatic differentiation, mostly python but some support for Java, C, Go)
- **Torch** – machine learning and scientific computing framework (Lua)
- **Scikit-learn** – library for general machine learning (python)
- **Caffe** – neural network framework (python interface, written in C++, popular in HEP, large library of pre-trained models)
- **Chainer** – neural network framework (python)
- **Lasagne** – neural network library over Theano (python)
- **Keras** – neural network library over Theano/Tensor Flow (python, also higher-level than Lasagne)
- **MATLAB** neural network toolbox

- Bare bones example of how things are structured in Theano and Lasagne

- Easy to set up mechanically

→ *much of the difficulty in using NNs comes with the training process and defining the initial problem correctly*

```
1 import theano
2 import theano.tensor as T
3 import lasagne
4
5 lin = lasagne.layers.InputLayer(shape=(None, 500), input_var=input_var)
6
7 l1 = lasagne.layers.DenseLayer(lin,
8     num_units=100, nonlinearity=lasagne.nonlinearities.tanh,
9     W=lasagne.init.GlorotUniform(gain=1), b=lasagne.init.Normal(std=0.001, mean=0.0))
10
11 l2 = lasagne.layers.DenseLayer(l1,
12     num_units=70, nonlinearity=lasagne.nonlinearities.tanh,
13     W=lasagne.init.GlorotUniform(gain=1), b=lasagne.init.Normal(std=0.001, mean=0.0))
14
15 l3 = lasagne.layers.DenseLayer(l2,
16     num_units=10, nonlinearity=lasagne.nonlinearities.tanh,
17     W=lasagne.init.GlorotUniform(gain=1), b=lasagne.init.Normal(std=0.001, mean=0.0))
18
19 out = lasagne.layers.DenseLayer(l3,
20     num_units=1, nonlinearity=lasagne.nonlinearities.linear,
21     W=lasagne.init.GlorotUniform(gain=1), b=lasagne.init.Normal(std=0.001, mean=0.0))
22
23 input_var = T.matrix('inputs', dtype='float32')
24 target_var = T.matrix('targets', dtype='float32')
25
26 prediction = lasagne.layers.get_output(out)
27
28 loss = lasagne.objectives.squared_error(prediction, target_var)
29 params = lasagne.layers.get_all_params(out, trainable=True)
30 updates = lasagne.updates.adam(loss, params, learning_rate=0.0001)
31 train_fn = theano.function([input_var, target_var], [loss, prediction])
32
33 #would then use the following to do one training update, where "inputs" and "targets"
34 #are your training data:
35 trn_loss, trn_pred = train_fn(inputs, targets)
```


Questions?

Backpropagation

Vectorized notation: $a_j = f(\sum_k w_{jk} x_k + b_j) \rightarrow f(wx + b)$

Layer-by layer: $a^l = f(w^l a^{l-1} + b^l) = f(z^l)$

a_j j^{th} node activation

f applied element-wise

b_j j^{th} node bias

$$\delta_j^l \equiv \frac{\partial C}{\partial z_j^l}$$

w_{jk} j^{th} node in layer l , k^{th} node in $l - 1$

$$\delta_j^{N_l} = \frac{\partial C}{\partial a_j^{N_l}} f'(z_j^{N_l}) \rightarrow \delta^{N_l} = \nabla_a C \odot f'(z^{N_l})$$

$$\delta_j^l = \sum_k \frac{\partial C}{\partial z_k^{l+1}} \frac{\partial z_k^{l+1}}{\partial z_j^l} = \sum_k \delta_k^{l+1} \frac{\partial z_k^{l+1}}{\partial z_j^l}$$

$$= \sum_k w_{kj}^{l+1} \delta_k^{l+1} f'(z_j^l)$$

$$\begin{aligned} z_k^{l+1} &= \sum_j w_{kj}^{l+1} a_j^l + b_k^{l+1} \\ &= \sum_j w_{kj}^{l+1} f(z_j^l) + b_k^{l+1} \end{aligned}$$

$$\frac{\partial z_k^{l+1}}{\partial z_j^l} = w_{kj}^{l+1} f'(z_j^l)$$

For each training instance:

1. Forward Pass:

For $l = 1, 2, 3 \dots N_l$

$$z^l = w^l a^{l-1} + b$$

$$a^l = f(z^l)$$

2. 'Error':

$$\delta^{N_l} = \nabla_a C \odot f'(z^{N_l})$$

3. Backward Pass:

For $l = N_l - 1, N_l - 2, \dots 1$

$$\delta^l = w^{l+1} \delta^{l+1} \odot f'(z^l)$$

4. Final Derivatives:

$$\frac{\partial C}{\partial w_{jk}^l} = a_k^{l-1} \delta_j^l \quad \frac{\partial C}{\partial b_j^l} = \delta_j^l$$